

AMENDMENTS TO THE CLAIMS

1. (Currently amended) A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the method comprising:

maintaining a ~~bitmap~~ record on the secondary storage subsystem, which is predictive of locations to which data are to be written on the primary storage subsystem by a host processor, the record including a designation of locations to which the host is expected to write in the near future;

receiving at the primary storage subsystem, from the host processor, the data to be written to a specified location on the first non-volatile storage media;

making a determination that the specified location is not included in the record, and responsively to the determination sending a message from the primary storage subsystem to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted ~~bits~~ locations in the record corresponding to the one or more further locations,

~~wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery;~~

signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, upon having made the determination that the specified location was not included in the record, responsively to receiving an acknowledgment at the primary storage subsystem from the

secondary storage subsystem indicating that the record has been updated; and

storing the data in the specified location on both the first and second non-volatile storage media.

2. (Original) The method according to claim 1, wherein sending the message comprises copying the data synchronously from the primary storage subsystem to the secondary storage subsystem.

3. (Previously presented) The method according to claim 2, wherein storing the data comprises, upon making the determination that the specified location is included in the record, copying the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without updating the record with respect to the specified location.

4. (Original) The method according to claim 3, wherein copying the data comprises transmitting the data between mutually-remote sites over a communication link between the sites.

5. (Previously presented) The method according to claim 3, wherein maintaining the record comprises maintaining a copy of the record on the primary storage subsystem, and wherein signaling the host processor comprises, upon making the determination that the specified location is included in the record, indicating to the host processor that the data have been stored without waiting to receive the acknowledgment from the secondary storage subsystem.

6. (Original) The method according to claim 1, wherein copying the data comprises creating a mirror on the secondary storage subsystem of the data received by the primary storage subsystem.

7. (Original) The method according to claim 6, and comprising, upon occurrence of a failure in the primary storage subsystem, configuring the secondary storage

subsystem to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

8. (Original) The method according to claim 6, and comprising, upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record, a portion of the data from the secondary storage subsystem to the primary storage subsystem for storage on the primary storage subsystem.

9. (Original) The method according to claim 1, wherein maintaining and updating the record comprise marking respective bits in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

10. (Currently amended) A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the method comprising:

maintaining a record on the secondary storage subsystem, which is predictive of locations to which data are to be written on the primary storage subsystem by a host processor, wherein maintaining the record comprises maintaining a copy of the record on the primary storage subsystem;

receiving at the primary storage subsystem, from the host processor, the data to be written to a specified location on the first non-volatile storage media;

making a determination that the specified location is not included in the record, and responsively to the determination sending a message from the primary storage subsystem to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record, wherein sending the message comprises deciding at the primary storage subsystem to send the message responsively to the copy of the record, and

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted ~~bits~~ locations in the record corresponding to the one or more further locations,

~~wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery;~~

signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, upon having made the determination that the specified location was not included in the record, responsively to receiving an acknowledgment at the primary storage subsystem from the secondary storage subsystem indicating that the record has been updated; and

storing the data in the specified location on both the first and second non-volatile storage media.

11. (Original) The method according to claim 10, wherein sending the message comprises modifying both the record and the copy of the record responsively to the specified location.

12. (Original) The method according to claim 11, wherein modifying both the record and the copy of the record comprises adding a plurality of locations, including the specified location, to both the record and the copy of the record.

13. (Original) The method according to claim 10, wherein maintaining the copy of the record comprises selecting one or more locations, other than the specified location, to be removed from the record, and instructing the secondary storage subsystem to remove the one or more

locations from the record, so as to limit a size of the record.

14. (Original) The method according to claim 13, wherein storing the data comprises copying the data to be stored in the one or more locations from the primary storage subsystem to the secondary storage subsystem, and wherein selecting the one or more locations comprises receiving a return message from the secondary storage subsystem indicating that the secondary storage subsystem has received the copied data, and selecting the one or more locations to be removed from the record responsively to the return message.

15. (Original) The method according to claim 13, wherein selecting the one or more locations comprises identifying the locations at which the first and second non-volatile storage media contain substantially identical data, and selecting for removal one of the identified locations that was least-recently added to the record.

16. (Original) The method according to claim 13, wherein sending the message comprises adding one or more entries to both the record and the copy of the record responsively to the specified location, and grouping the entries added to the copy of the record and the record in generations according to an order of adding the entries to the records, and wherein selecting the one or more locations comprises determining at the primary subsystem that all the entries in one of the generations may be removed from the record.

17. (Original) The method according to claim 13, wherein instructing the secondary storage subsystem to remove the one or more locations comprises appending an instruction to the message sent from the primary storage subsystem to the secondary storage subsystem.

18. (Canceled)

19. (Previously presented) The method according to claim 1, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

20. (Previously presented) The method according to claim 1, wherein maintaining the record comprises recording the locations to which the data are written using an object-based storage technique, and wherein the one or more further locations are chosen based on a logical connection between storage objects.

21-60. (Canceled)